

Maintenance Plan

By E. J. AITCHESON and R. F. COOK

(Manuscript received July 15, 1970)

The No. 1 ESS ADF message switching system provides a store and forward data service which places special demands on system dependability and maintainability. This paper discusses the hardware and software features used to detect and sectionalize troubles, as well as the recovery techniques used to restore service quickly. Maintenance of the line facilities, use of circuit redundancy, and message data protection are also included.

I. INTRODUCTION

A communication switching system must be designed with dependability and maintainability as an integral part of the overall plan. The No. 1 ESS ADF store and forward message switching system is no exception. Continuous high quality service is of vital importance. The characteristics of high quality data service include good error performance, 24-hour service with a minimum of interruptions, fast restoration of service, and no loss of messages when interruptions do occur.

The system's error performance objective for basic station-to-station messages is: on the average, no more than one error in 10^5 bits—99 percent of the time while continually transmitting. The error performance will be determined largely by the station access lines since the error rates within the switching office are much lower. The switching center hardware was designed to include optional error detection and correction features (by retransmission) to achieve even greater transmission accuracy.

The reliability objective for the No. 1 ESS ADF system is to provide continuous service with system downtime not exceeding 2 hours in 40 years. The store and forward data features make it possible to preserve message information under the most severe fault conditions so messages can be retransmitted or retrieved when service resumes. Once a store and forward message office accepts incoming traffic for

delivery at a later time, it is of utmost importance that the message and delivery stimulus are never lost.

The No. 1 ESS ADF maintainability objectives provide a system whereby most faults can be located automatically and repaired quickly with minimum effect on service.

II. GENERAL MAINTENANCE PLAN

In the No. 1 ESS ADF system, all message data is routed through common processing units. The transmitted teletypewriter data from user stations is converted into computer words by an autonomous data scanner-distributor (DSD) and the autonomous buffer control. The message is assembled into information blocks in call store memory, buffered for delivery in a disk memory, and permanently stored on magnetic tape for retrieval purposes. The consequences of a failure in these common units, through which all messages may pass, can be severe. Fast recovery from failures is vital, as interruptions can affect all messages in the process of being transmitted or received. For example, if buffer control I/O processing is interrupted for longer than 66 milliseconds, input messages from all 150 baud stations must be retransmitted. To avoid complete system failure when a single component fails, circuit redundancy is used. With circuit redundancy, service can be maintained during fault diagnosis, fault repair, and routine maintenance.

The maintenance goal is to recover from faults before service is appreciably affected so that the user is unaware of trouble. To accomplish this goal, errors and faults must be detected quickly before incorrect information propagates into other units in the system. Continuous hardware checks provide the principal means for detecting faults in the common processing units. When a hardware check fails, an interrupt sequencer in the central control transfers program control to maintenance fault recognition programs. These programs isolate the faulty unit and switch a duplicate unit into service. The standby duplicated units are normally run in synchronism with the active unit to keep the contents of standby units up to date, thus making them instantly available for use when a faulty unit is removed from service. For many faults, the trouble detection and reconfiguration process is sufficiently fast to avoid service interruptions from a user's viewpoint.

When fault recognition programs experience difficulty in restoring service, error analysis routines are used. Error analysis programs record a history of system interrupts, troubles, and configurations. These

programs are used in conjunction with fault recognition routines to isolate units with marginal faults or with faults that are difficult to locate. The error analysis programs, which use a statistical approach to fault isolation, can be considered as a backup to assist in recovering the system.

Interruptions in service may occur for some faults that are difficult to locate. In these cases, the customer automatically receives service messages that will assist in determining the corrective action to be taken. Interrupted input messages must be resubmitted for delivery to the office by the user. Interrupted output messages will be retransmitted to the station automatically under program control.

After call processing has resumed, diagnostic programs are scheduled to be run on the unit removed from service. The purpose of these programs is to test the unit thoroughly and to supply test results to the maintenance craftsman. Maintenance trouble locating manuals translate the test results and list the circuit packs that might be faulty.

The system also includes fault detection capability for facilities dedicated to a user's line. Automatic in-service performance checks executed by the system are used to test both active and idle lines. Troubles that degrade user service can be detected and corrected before they become catastrophic; for example, parity over each character in the message checks terminal circuits and the quality of the transmission facility. Failure of the station to respond correctly to polling signals sent by the switching center can initiate corrective action for idle lines. When line faults or marginal station troubles are detected, the system is not interrupted. A control serving test center is notified of the problem by a teletypewriter message, where the necessary action is taken to sectionalize and clear the trouble.

The user is also provided with service features, which can be used when difficulties are encountered. For example, the user can request retransmission or retrieval of messages that were received with errors. Alternate terminals can be specified to receive messages addressed to a faulty terminal. Traffic statistics can be requested periodically that include the number of messages delivered to and transmitted by each station.

The dependability of the system is enhanced by the use of conservative circuit designs and long-life components. Wherever possible, tried and proven No. 1 ESS units, packs, and components are used. The same design principles, using liberal operating margins, worst-case circuit designs, long-life silicon and magnetic devices that have

proved effective in past projects are applied to this message switching system to obtain reliable units and a low trouble rate.

The principal features of the maintenance plan are as follows:

(i) Conservative circuit designs and long-life components are used to obtain reliable units.

(ii) Redundant units are used to provide service in the presence of failures and for routine preventive maintenance.

(iii) Rapid detection of faults by continuous hardware and software checks.

(iv) Recovery procedures by fault recognition programs are designed to preserve message information while testing and configuring the system around faulty units.

(v) Error analysis programs are used to distinguish between occasional errors and marginal or intermittent faults.

(vi) Diagnostic programs, interleaved with message processing programs, are automatically scheduled to isolate faults to replaceable plug-in circuit packages.

(vii) In-service checks of user lines provide rapid detection of faults and marginal troubles.

The following sections describe the redundancy plan, maintenance circuits, and maintenance programs. Those maintenance features for the central processor and other No. 1 ESS units are covered in the No. 1 Electronic Switching System described in the September, 1964, issue of the B.S.T.J.

III. CIRCUIT REDUNDANCY

The ADF system consists of a No. 1 ESS central processor and a community of ADF units to perform the store and forward message switching functions (Fig. 1). These units include an autonomous data scanner-distributor¹ to access the lines, a message store² (disk store) to assemble and hold messages awaiting delivery, a magnetic tape store³ to provide a permanent file for messages, a buffer store⁴ for scratch pad use, and a buffer control to perform repetitive tasks related to disk, tape, and I/O operations. Operational programs load commands and data for the buffer control into dedicated task dispenser queues in the buffer store. The queues are unloaded by independent wired logic sequencers in the buffer control which interpret the commands and perform the data transfers.

As shown on Fig. 1, the buffer control, buffer stores, message stores,

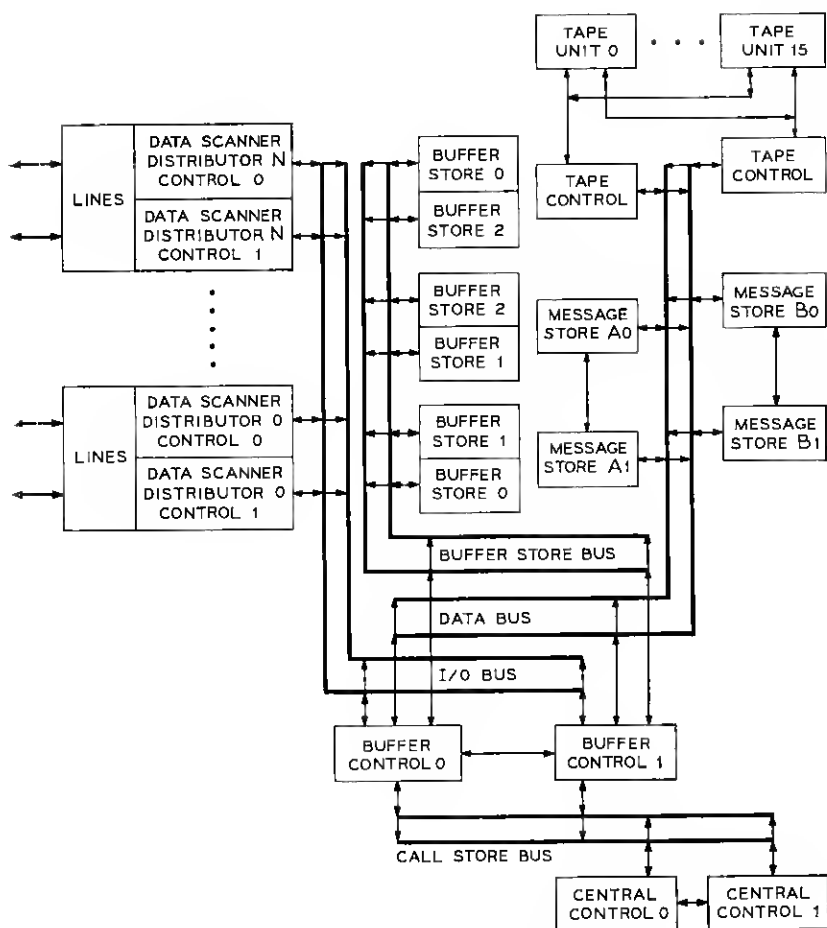


Fig. 1—Duplication of message processing units.

and communication buses are duplicated and the units are operated in a synchronous matching mode.

The autonomous data scanner-distributor units are partially duplicated. These units are used to convert input message characters, that arrive as a serial bit stream, into characters which are sent to the buffer control in parallel word form over the I/O bus. For output messages, the autonomous data scanner-distributor units receive characters from the buffer control in parallel, convert these characters into a serial bit stream, and route the data to the designated output line

terminal. The line terminal logic in the autonomous data scanner-distributor is not duplicated since faults in this logic will affect, at most, only 8 lines. The remaining logic in the autonomous data scanner-distributor unit, which uses time division techniques to perform the serial-to-parallel conversion and buffer the data, is duplicated.

Two tape unit controls operate in a simplex mode to provide simultaneous, but independent, tape storage operations. Under fault conditions, this redundancy allows messages to be put on a permanent tape file while deferrable tasks, such as message retrieval, are postponed. Each tape unit control can access a maximum of 16 tape units that provide sufficient spares for normal tape mounting and demounting operations, as well as for routine maintenance.

All buffer control data communication buses are fully duplicated. Each unit can be configured to receive data from either bus or send data on either or both buses by means of route control flip-flops. Normally each unit is configured to send and receive data on the same bus. Half of the duplicated units or controllers send and receive data on bus 0, while the other half use bus 1. When one unit is removed from service for maintenance purposes, the routing for the other unit is configured to retain as much of the duplicated system as possible. If the remaining unit sends data on both buses, then the buffer control can continue to access and match all other duplicated units on the same bus system in a normal manner. Because the tape unit control is not duplicated, it receives data on one bus and sends data on both buses to the buffer controls.

IV. BUFFER CONTROL COMMUNITY MAINTENANCE

The buffer control coordinates the transmission of data between all ADF units and verifies that these units and buses are functioning correctly. A malfunction in an ADF unit may be discovered by buffer control through several sources, which include parity failures during a bus transmission, status reports from the units, match failures at the buffer control, or by the failure of a unit to send buffer control an all-seems-well (ASW) response. The buffer control may react to these malfunctions by repeating the failed operation, incrementing error counters, reporting the trouble to operational programs via software queues, by interrupting normal processing with a maintenance interrupt, or by a combination of the above actions. The circuit features used to detect and report troubles in the ADF units are covered in the following sections.

4.1 *Interrupts*

When troubles are detected in the system, a wired sequencer in the central control interrupts the program in progress and transfers to a maintenance program that determines the source of the interrupt and takes corrective action. It is possible that more than one unit may detect and report a fault at the same time. To handle this problem, the trouble sources are grouped into interrupt levels and ranked according to the seriousness of the trouble source. Interrupt levels A through E are caused by central processor related faults or are manually induced.⁵

The ADF equipment, consisting of the buffer control and its peripheral communities, can generate F-level maintenance interrupts when malfunctions are detected. The interrupt will always be issued by the active buffer control, which is the only ADF unit that can interrupt the central control directly. An ADF peripheral unit can cause a maintenance interrupt only by inhibiting its all-seems-well signal to buffer control. This, in turn, will cause the buffer control to issue the maintenance interrupt to central control with only a single functional sequencer stopped.

If the central control receives an F-level maintenance interrupt, it will transfer program control to the F-level filter program. If the F-level source is the buffer control, the filter program will interrogate the buffer control error indicators to determine which buffer control or peripheral community is at fault. Failures of the central control peripheral units are also sources of F-level interrupts. Once the source is determined, the filter threads together the fault recognition programs to be executed to isolate and configure around the faulty unit.

Software checks of buffer control operations can detect errors and transfer control to maintenance routines. Since the central control and buffer control communicate with one another by software task dispensers, the main program can detect functional troubles by monitoring the progress and status of these queues. When buffer control completes a task, it overwrites the command in the queue with a passing or failing status report. If an operational program detects that the queue contains incorrect status reports, it can enter J-level fault recognition routines to test associated hardware. The fault recognition routine reissues the command on a half directed basis. In this manner, the faulty unit, not able to process the command correctly, is isolated.

4.2 *Fault Detection Maintenance Features*

4.2.1 *All-Seems-Well*

Each time the buffer control addresses an ADF peripheral unit, a 1-bit ASW signal from the unit is expected. The ASW signal indicates that the maintenance checks performed on the bus instruction have passed and the unit is functioning correctly. The ADF unit informs the buffer control that an error has been detected by inhibiting its ASW signal.

Table I summarizes the maintenance checks performed by the ADF peripheral units causing ASW failures. The disk and tape unit control perform other checks not shown in Table I that are reported through the use of an instruction queue as described in Section 4.1.

The buffer control reacts to the ASW failure as follows:

(i) The buffer control will first repeat the instruction and cause an F-level interrupt only if the repeat fails. On the other hand, if the ASW fails on a class of instructions referred to as central control read instructions, then the ASW failure is passed on to the central control by inhibiting the ASW on the call store bus. For this case, the central control has the responsibility of repeating or interrupting the system.

(ii) The ASW failure also selectively stops the logic sequencer in the buffer control responsible for the bus instruction that failed. By doing this, the state of the logic is frozen, thereby preventing the fault from

TABLE I—MAINTENANCE ALL-SEEMS-WELL CHECKS

Message Store

1. Parity on instructions received from buffer control.
2. Synchronization of duplicated disks (servo check).
3. Internal clock check.
4. Maintenance order received during normal operation (mode check).
5. Buffer control out of sync with disk (sector match check).

Tape Unit Control

1. Parity on data received from buffer control.

Autonomous Data Scanner-Distributor

1. Match of input data from lines.
 2. Match of output data to lines.
 3. Match of data sent to buffer control.
 4. Match of time slot address.
 5. Parity from input line to bus access.
 6. Parity from bus access to output line.
 7. Parity over bus address and data.
 8. Address translator.
 9. Buffer control fails to acknowledge data sent by the DSD.
 10. Unit name decoder for bus instruction.
-

propagating. Other sequencers in the buffer control are allowed to continue normal processing until the fault recognition routines enter to test the buffer control.

4.2.2 *Parity Protection*

Data transmitted between the ADF units is protected by parity bits. The ADF peripheral units check parity on instructions received from the buffer control and will inhibit the ASW if the parity check fails. The buffer control checks parity on all data received from the peripheral units with one exception: central control may read from memory locations or registers in any ADF unit. The data from these central control read instructions is passed through buffer control, and central control is responsible for checking parity and reacting to parity failures. Most parity failures are treated similar to ASW failures. However, a special block repeat procedure is used for parity failures on instructions which read data from disk. A parity failure on disk reads is recorded for later use and buffer control finishes reading the block of data from disk. The buffer control sets a repeat flag in the queue status word and the buffer control rereads the entire block at a later time. If the block repeat fails, a program which administers the disk instruction queue calls in a fault recognition maintenance program. F-level interrupts do not occur for this type of failure.

4.2.3 *Error Rate Check*

Each bus sequencer is designed to automatically retry an operation if a parity or ASW failure is detected on the first attempt. A match interrupt is inhibited by an ASW or parity failure on the first attempt. Each time a bus error (failure on first try, success on second try) is encountered, an error counter dedicated to the bus is advanced by 1. When a count of 32 is reached, an overflow bit is set. This bit is periodically scanned and the error counter reset by the central control under command of a maintenance program. If the overflow bit is set as a result of an error rate in excess of a predetermined software threshold, the bus sequencer is forced to stop on the first failure and generate an F-level interrupt. In this way, the unit causing a high single error rate can be identified and removed for diagnosis.

4.2.4 *Matching*

The buffer control community and its buses are fully duplicated and run in a synchronous matching mode. All external bus operations are

matched, bit by bit, using hardware matchers. The information sent on the buses includes an address field, used to access a specified register within a unit, and a data field. The active buffer control matches address and the standby matches data. A mismatch in either buffer control will cause the bus sequencer in both buffer controls, handling the operation, to stop and freeze the bus priority F/F associated with the sequencer using the bus during that cycle. An F-level maintenance interrupt is then sent to both central controls by the active buffer control. Operations by other sequencers not requiring the stopped bus are unaffected. Normally, match failures cause an immediate interrupt. However, if parity or ASW failures also occur at the same time, then the instruction may be repeated as described in Sections 4.2.1 and 4.2.2.

A directed or off-normal match mode is provided where the circuits to be matched and the time the match is to take place are specified by program. This mode is used by the buffer control diagnostic program.

4.2.5 *Internal Sequencer Check*

All internal wired logic sequencers in buffer control are designed to advance through a wired series of sequencer states and, upon completion, recycle to a starting point. These sequencers receive their external stimulus from associated peripheral controllers in the form of service requests. Internal stimulus is provided by permission to use an internal or external bus. The response to this stimulus is controlled by the hard-wired sequencer logic. Since at each point the sequencer knows what to expect next, wired checks are made to verify the sequencer operation. Thus, invalid or out-of-sequence service requests (external stimulus) can be detected by the buffer control. They will cause that sequencer to stop and generate an F-level maintenance interrupt. All sequencer faults not detected in the above manner will be detected when that sequencer attempts to perform an external bus operation. If the duplicated sequencers are out of step, a data and address mismatch will result when one sequencer attempts to use the bus and its mate does not. Should the fault occur in an area common to the buffer controls, such as the service request decoders, the external peripheral equipment being addressed will inhibit its ASW response because it receives an out-of-sequence order.

4.2.6 *Clock Checks*

The buffer control clock is a 22-phase ring feedback chain driven by a 2-MHz source provided by the active central control. The phase

relationship of the ring is synchronized to that of the central control every $5.5 \mu\text{s}$ using a sync pulse generated by the active central control. Checks are made to verify that the 2-MHz clock is present and that the clock's phases are generated correctly. A clock fault stops all sequencers in both buffer controls and generates an immediate F-level interrupt. Error indicators related to clock circuits are accessed by scan points external to the buffer control circuits so the fault recognition and diagnostic programs can isolate the faulty unit without requiring an internal buffer control bus read.

V. MAINTENANCE PROGRAMS

5.1 *Fault Recognition and Recovery*

Fault recognition programs are called in when errors or faults are reported by the interrupt logic, base level, or low priority nondeferable programs.⁵ The purpose of the fault recognition programs is to determine the source of the trouble, remove the faulty unit from service, and restore the system message processing capability as quickly as possible. These programs also distinguish between errors and faults and may take no action other than recording that an error was detected. After corrective action has been taken, recovery routines initialize hardware and return to normal processing as gracefully as possible. In many cases, processing resumes at the point where the interrupt occurred. The fault recognition and recovery process emphasizes fast recovery to avoid destroying message information. In addition, special procedures are used to insure that messages on disk are not destroyed when severe problems are encountered.

5.1.1 *Disk Recovery and Message Protection*

There are two duplicated disk communities, each capable of storing 57 million bits of binary information. A portion of this data provides a present and past history record for the entire system and must be maintained over long periods of time. If one disk gets out of date and its mate experiences a failure, the system loses its ability to retrieve from that community. An out of date disk is never automatically configured into the system. Two separate recovery strategies are provided for this situation. The first requires a manual emergency action phase 5, which will clear all past history (time 0 start) and bootstraps the switcher into a workable configuration. Because the No. 1 ESS ADF is a store-and-forward system, hundreds of messages awaiting delivery would be permanently lost and no notification could

be sent to the sender to retransmit undelivered traffic. To avoid this gross loss of traffic, a second recovery plan is provided. When the duplex disk failure is encountered, all operational processing is halted, and notification of a duplex disk failure is given at the maintenance control center. Office personnel can then examine maintenance teletypewriter printouts and decide which disk file had the last active copy of system records. The maintenance craftsman then protects the file from being overwritten by simply retracting the read/write data heads away from the memory surface. When the disk controller trouble is cleared, the protected disk is bootstrapped into the active system by a manual emergency action phase 4 restart initiated at the master control center. A phase 4 restart bootstraps the equipment and initializes the call stores and buffer store communities. Disk records are assumed to be accurate. All traffic being held for delivery at the time of the failure is then delivered in a normal fashion.

5.1.2 *Buffer Control Recovery*

The buffer control contains a number of sequencers that must be initialized before the buffer control can be restored to service. For example, the disk sequencers in buffer control must be synchronized with the disk and always know the address (sector) positioned under the reading and writing heads. The buffer control uses three types of service request signals from the disk to aid in the communication between these two units. The disk is divided into 16 pie-shaped sectors. At the start of each sector, the buffer control receives an instruction request. The buffer control responds with an instruction, telling the disk the operation to be performed during the sector, as well as the specific data location addresses involved in the instruction. While the disk is moving through the sector, the data is transferred between the buffer control and the disk in response to data request signals sent from the disk. At the end of the sector, the disk sends a status request that signals the end of the operation. The buffer control reads the status report from the disk which indicates the present sector address and contains trouble status information. The status information is loaded into an instruction queue and examined at a later time by program. Before a buffer control can be restored to service, the disk sequencer in the buffer control must be initialized with the present sector address.

To achieve synchronization, the buffer control disk sequencer, under program control, is initialized to look for status requests only. When

coincident status requests are received from the two duplicated disks, the sequencer reads the disk status reports. In the start up mode, the sequencer extracts the four bits from the status report which corresponds to the current disk sector being accessed by the disk controller. The sequencer queue counter is set to the value of these four bits plus 1 (+1). The sequencer then resets the start up control flip-flop, and advances to state 0 to preload a task for the next sector. Thus, the next instruction request received by the buffer control is honored and the test to be performed is executed.

The tape sequencer must be initialized in one of two states, depending on what it was executing when the stop occurred. If a buffer control to tape transfer was in progress, the sequencer must be initialized to honor a status request. Otherwise, it is initialized to look for a new instruction request. The queue counter must be readjusted since it acts as an operational job pointer. Since the queue may contain an operational tape stop operation, the queue counter must be set up so all tasks will be executed before reaching the operational stop.

Each of these decisions and the initial state of the sequencers are set up by a software maintenance quickstart program. Once the hardware is initialized, startup is directly associated with external hardware stimulus. Restart is only required after buffer control has been stopped by a fault or a maintenance program.

5.1.3 *Error Analysis*

The fault recognition programs are designed to restore a faulty system to normal operation within a few milliseconds. The fault recognition programs accomplish this objective for most faults. However, these programs do not have time to exhaustively test suspect units because message handling will be affected each time a maintenance interrupt occurs. (The fault recognition programs must make a decision on the basis of a brief examination of the suspect units. For most faults, the correct unit is removed from active service and processing continue without any loss of data or service.) Some marginal faults are more difficult to isolate and fault recognition may not discover the fault or may remove the wrong unit from service. Maintenance interrupts will continue to occur until the faulty unit is isolated from service. Error analysis and emergency action routines are used to restore service when persistent interrupts occur. The error analysis routines keep a record of error counts, previous system configurations, and the active-standby status of the units to assist the fault recogni-

tion programs. If interrupts continue to occur, more drastic action is taken by emergency action routines.⁵

5.1.4 *Monitor Mode*

An electromechanical device, such as a disk file, can generate a low level of errors which are not reproducible during fault recognition testing. Although the situation must be ultimately corrected, its minimal effects on the operating system warranted taking several seconds to allow careful programmed analysis of the trouble condition. The most serious consideration is to avoid removing the wrong disk from the active system, thereby causing its contents to get out of date with the active copy. The updating process requires about six minutes and assumes read access to the entire active disk. The fault recognition monitor program interrogates software error counters to detect if a predetermined error threshold has been exceeded by a duplex disk community. Once the error rate has been exceeded, and its source not isolated to a suspect disk, the fault recognition program selects a disk and configures it to respond for both itself and its mate. Rather than removing the remaining disk completely from the active system, it is configured to listen and record only. If the errors persist, the unit removed can be immediately restored to active service. (No update required.) Its mate is then assumed to be faulty and can be completely removed for programmed diagnosis. Although simple, this technique has been extremely effective in preventing user messages from being permanently lost prior to delivery.

5.1.5 *Bootstrap Routines*

Emergency action recovery of ADF equipment is accomplished by software bootstrap programs. This maintenance software decouples all bus configurations and rejoins simplex equipments in a semirandom fashion. Once a complete system is established, it is restarted and monitored for excessive interrupts over a short interval. If interrupts continue, the bootstrap software is repeatedly entered. Since this process is semirandom, a working system will be established, possibly after multiple attempts. The faulty unit is detected when an automatic diagnostic program is executed before joining it to the already working system. Units passing the diagnostic program are updated and joined to the working half.

To avoid simplexing and having mass memory become outdated, two types of bootstrap routines are employed. A hard bootstrap sim-

plexes all units and rejoins them only after a successful diagnostic test. A soft bootstrap which takes less time configures units according to their last known status record that is maintained in the call store complex. Both hard and soft type bootstraps are threaded together for a particular recovery strategy.

5.2 *Diagnostics*

The purpose of diagnostic programs is to thoroughly test a unit that has been removed from service and to generate sufficient test data to isolate the fault to within a few replaceable circuit packs. These programs are run in short segments interleaved with call programs. The diagnostic program for a unit consists of a control program and a series of test routines that are followed in a fixed sequence. These test routines are grouped together to form a sequence (phase) which tests a specific function in the unit, such as the bus access logic. The buffer control, for example, has a diagnostic consisting of a control program and 28 phases of tests.

When a diagnostic test fails, two courses of action are possible. The remaining tests may be run to get additional test data, thereby more accurately pinpointing the faulty pack; or, the diagnostic may be terminated on the basis that further tests will generate inconsistent or misleading results. For either case, the maintenance teletypewriter printout displays the phases that failed, the test results in an octal code (raw data), and a 12-digit trouble number. A maintenance dictionary is used to translate the trouble number by listing all circuit pack locations that could cause the trouble number. (Figure 2 shows a typical teletypewriter trouble report for a fault in tape unit control zero.) The diagnostic results are listed, which include the universal trouble number. A section of the tape unit control dictionary is also shown in Fig. 2 listing the pack location for that trouble number. If the trouble number generated by the diagnostic is not found in the dictionary, then the raw test data listed in the teletypewriter printout is analyzed. A manual which lists the tests, the expected test results, and the logic circuit tested is available to aid in resolving marginal or inconsistent faults.

5.3 *Routine Exercise*

All units in the system are periodically removed from service and diagnosed to test maintenance error detection hardware. This insures the ability of a unit to detect and respond to faults in operational

sequencers. If no source for a stopped sequencer can be determined, the sequencer is reinitialized and restarted.

Critical parameters and constants are stored in the buffer stores. These constants are related to the number of autonomous data scanner-distributors the buffer control should scan: class of service, character type, speed, and error control. Should the data become overwritten or otherwise destroyed, message processing can stop. To prevent this critical data from being lost and going unnoticed for long periods of time, the system audits the area every 8 seconds. If a bad data word is found, the audit will initiate an emergency action phase 1, causing all of the critical constants stored in buffer store to be reinitialized.

Under normal operating conditions, no maintenance interrupts, the operational processing program is the only means by which the system can be alerted of trouble conditions. These programs, in addition to performing their operational work, must verify that data is kept moving to the various peripheral controllers. Normally, faults which cause buffer control to stop processing disk data also cause a maintenance interrupt which brings in fault recognition and diagnostic programs. However, under unusual circumstances, a maintenance program which has temporarily inhibited disk service request signals may be aborted. Under these conditions, buffer control sequencers are stopped because of lack of stimulus from the disk. The hardware audit would find the sequencer stop flip-flop reset (normal) and release control back to normal processing. For this class of fault, the operational program administering the task queue must schedule a base-level fault recognition test. This program will first interrogate the buffer control error indicators and find no flags set. It will then thread-in a software buffer control restart program so it can monitor the queue. Restarting the buffer control will cause the inhibit service request flip-flop to be reset and the sequencer will cause the backed up task in the queue to be executed. The fault recognition program will conclude all is normal and release control. The operational program, detecting that the queues are now being processed, will discontinue entering maintenance routines. Other sequences are protected from being left stopped in the absence of maintenance interrupts in a similar manner.

VI. AIDS TO MANUAL PROCEDURES

Although most of the switching center maintenance procedures are accomplished using direct program control, other semiautomatic and

manual test procedures are provided where they have become necessary.

Off-line equipment tests can be executed by configuring a central control and the desired equipments on one-half of the duplicated bus system, while the active system is running normally on the other. The active central control can then be made to start and stop the off-line central control, causing it to execute program orders. Normal call processing is unaffected.

Each time program control enters base-level work, the software operates a central pulse distributor point. This pulse drives a meter calibrated in milliseconds, indicating call processing activity. When maintenance software is being entered excessively by soft or hard interrupts or the system is operating with a unit causing high rates of single errors, the meter will fluctuate and indicate higher values. This meter provides a continuous indication of traffic load at the maintenance control center, and alerts office personnel of a trouble condition that must be closely monitored.

An ADF office is equipped with special equipment to maintain disk files. A special disk exercise unit is provided so that all disk addresses can be tested off-line. This test set is used to verify a disk when returned from the factory after repair. In addition, special cleaning and purging equipment is on hand to maintain disk files. A master disk clock writer is also provided to write the clock (program) onto a new file received from the factory.

VII. MAINTENANCE DICTIONARY PRODUCTION

The maintenance dictionaries are used to convert the diagnostic results received from the maintenance teletypewriter to a list of circuit pack locations. The dictionaries were generated basically in the same way as the No. 1 ESS dictionaries—that is, by inserting faults in a test model unit, running the diagnostics, recording the diagnostic results, sorting the data, and printing the results along with the package location. However, improvements were made in the fault insertion procedure. A program was written to search a Western Electric tape containing wiring information for all circuits in a unit. The information was used to generate a complete list of faults for all circuit packs. Faults for spare or unused circuits on these packs have been automatically eliminated. The list of faults for each pack was coded on punched cards and used to control the fault insertion equipment. Programs were designed to store the diagnostic results on any specified

No. 1 ESS ADF permanent file 9-track tape. An IBM computer was used to process the tapes, compute trouble numbers, and print the dictionaries. Approximately 200,000 faults were inserted to produce the dictionaries for the No. 1 ESS ADF units.

The results of the dictionary production indicate the diagnostics are about 85 percent effective in locating the faults detected by the maintenance hardware and software checks. About 15 percent of the faults inserted were not located to replaceable units by the diagnostics or produced inconsistent results. Manual procedures must be used in conjunction with off-line operation to isolate faults not detected by the diagnostics.

To produce a more effective diagnostic and dictionary, better and faster feedback is needed than is possible in the dictionary manual fault insertion procedures. The use of large scale digital computer simulation of logic circuits appears to be the answer.⁶ By using digital simulation, diagnostic design and fault insertion can more nearly parallel the logic design phase of a system. Program and logic changes can be made to isolate nearly 100 percent of the faults simulated before hardware designs are frozen. It is likely that future designs will depend heavily on digital simulation to produce fault dictionaries, and significant improvements in the effectiveness of diagnostics can be expected.

VIII. LINE MAINTENANCE

8.1 *Introduction*

The purpose of the line facilities is to provide an interface between the user's station and the common processing units. The ADF maintenance plan includes features for detecting, reporting, and isolating troubles in this equipment dedicated to a user's line. In general, the line maintenance approach is somewhat different from the plan for maintaining common hardware. For example, fault detection in common hardware is based mostly on hardware checks, and the system is interrupted when faults are detected. On the other hand, faults in line facilities are detected by both hardware and software (mostly software), and the system is not interrupted when faults are detected. Message processing continues while fault isolation tests are made. The maintenance procedures involve facilities at the customer location, at local or remote test centers, and at the switching center. Line maintenance, therefore, requires the cooperation of craftsmen at several locations and involves manual as well as programmed tests.

8.2 Test Center

The data lines from user stations have appearances at test boards in serving test centers. A line from the switching center to a user terminal could pass through several test centers (as shown in Fig. 3). The serving test center closest to the switching center, through which all lines must pass, is called the control serving test center. The craftsmen at the control serving test center are responsible for troubleshooting and maintaining the transmission and terminal facilities. Troubles with subscriber lines may be detected by monitoring circuits within the test center or by monitoring circuits and in-service message tests at the switching center. When line troubles are detected by the switching center, the test center receives service messages, indicating the type of check that failed and the identity of the station or line in trouble. All test centers involved in suspect

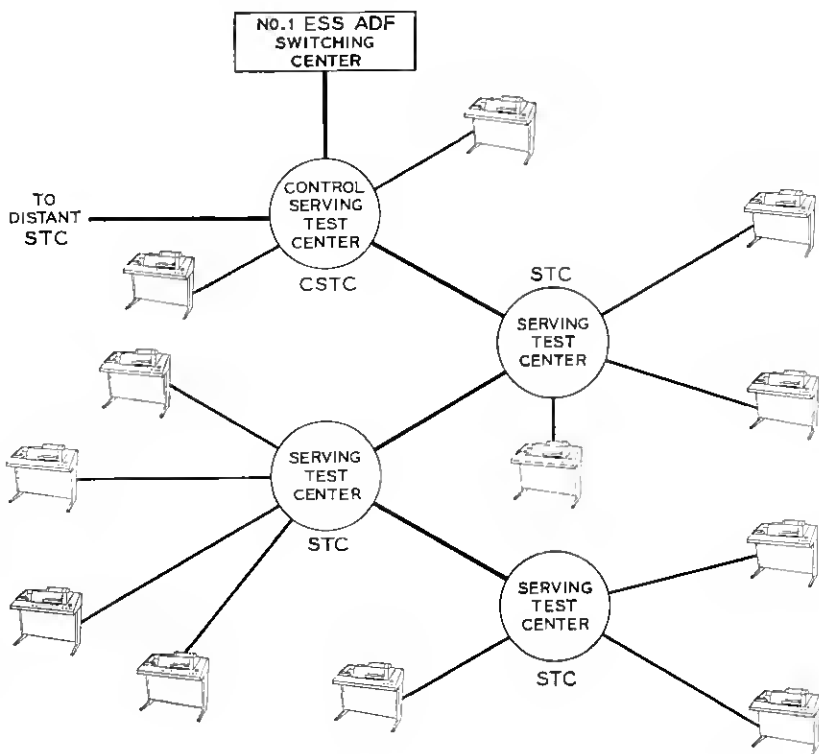


Fig. 3—Facilities for testing user lines.

lines cooperate in testing the complete facility. As part of the troubleshooting procedure, the control serving test center may send action request messages to the switching center requesting certain tests be performed. These switching center tests do not interrupt service and can detect marginal conditions that degrade service as well as facility failures.

8.2.1 *Control Serving Test Center Features*

The main facilities provided at the test center include a test service board of the type used for private line telegraph service, monitor teletypewriters, and a teletypewriter station to the switching center that is serviced in the same way as a user station. These facilities are used to perform tests that include the following:

- (i) A continuous open-line monitor checks the line for breaks.
- (ii) A high signal distortion check monitors the quality of the signals.
- (iii) Loop tests from the test center test board to the data set at the switching center and back to the test board are used to check the link between the test center and the switching center. This test requires a special circuit pack be inserted in the data set at the switching center to connect the send and receive lines.
- (iv) Loop tests from the test board to the user's terminal and back to the test center can sectionalize faults in the transmission facilities, station controllers, and terminals.
- (v) A monitor teletypewriter is available to manually test stations by sending character sequences or to monitor the line.
- (vi) A patching capability is available to transfer user facilities to spare lines between the test center and the switching center when transmission or terminal troubles are encountered.

8.2.2 *Station Facilities and Action Requests*

The test center must work closely with the switching center and make full use of the system's capability. The switching center can detect service degradation, line troubles, and assist in the test procedures. To communicate with the switching center, teletypewriter stations are used to receive service messages and to send action requests. These stations are connected to the switching center through the autonomous data scanner distributor and are serviced in the same way as a user station. Requests may be sent by the test center per-

sonnel requesting the switching system or craftsman to perform the following functions:

- (i) Place a station on skip, intercept, or alternate delivery.
- (ii) Change a station from one autonomous data scanner-distributor port to another. This action may be prompted because of troubles in the facilities between the test center and the switching center or an autonomous data scanner-distributor.
- (iii) Restore a station to its assigned line.
- (iv) Perform a distortion measurement on a specified input line.
- (v) Provide a status report of each station on a specified line, such as stations which are on alternate delivery, hold, or skip.
- (vi) Stop the delivery of messages to a station or cause a station to stop sending a message.

8.3 *Switching Center Line Maintenance*

8.3.1 *In-Service Checks and Service Messages*

The No. 1 ESS ADF office is programmed to perform in-service checks on messages being processed. For example, the switching center sends special characters to determine the status of the stations, to prepare the stations to send or receive messages, to terminate messages, or to verify reception of messages. If the station fails to respond or gives an incorrect response, the office repeats the sequence; if the failure repeats, the test center is informed of the trouble. The craftsmen at the test center are actively working on clearing troubles before the user recognizes a trouble. Service messages informing the control serving test center of line troubles include:

(i) *Polling failure*—Idle stations are periodically polled to determine if the stations have input messages. Failure to receive a valid "no message" response or a "yes, I have a message" response is reported. The polling procedure, therefore, provides a continuous check on the ability of idle stations to communicate with the office.

(ii) *Transmitting call enquiry code failure*—A polling response may indicate that a station requests to send a message. In this case, part of the message origination procedure requires the office to send a call enquiry code (CEC) to the station. If the station responds correctly with a start of heading (SOH) character, the office then sends heading and message number information to the station. A failure to respond correctly to the CEC is reported.

(iii) *Failure to restart*—After the office has sent heading and message number information to the station, the office restarts the station

teletypewriter transmitter. The station then sends the message. A failure to respond to the restart is reported.

(iv) *Station call-in failure*—Before a message is delivered to a station, the office determines if the station is ready to receive. Failure of the station to send a valid "ready" or "not ready" response is reported.

(v) *Roll call failure*—After a message is delivered to one or more stations on a line, the office roll calls each station receiving a message. If a negative roll call response is received, the message delivery is repeated. If a negative roll call response occurs again, the trouble is reported.

(vi) *Loss of control*—The office reports it has lost control of a line when a transmitting station will not respond to a request from the office to stop sending.

(vii) *Loss of facility*—The line from the test center terminates in a data set at the switching center. An open line will be detected at the data set by a ferrod monitor. A failure will be reported if an open line exists.

8.3.2 Character Parity

For stations that use the ASCII code, a parity bit is included for each character. The buffer control checks the parity of each character and replaces the character with a slash (\) symbol if the parity fails. When the message is delivered, the customer has the option of requesting the message originator to retransmit the message if a vital character was lost. The terminating station also checks parity on each character, and an underline (_) character is printed if a parity failure is detected. In this case, the error occurred in the facilities used for message delivery. If vital characters were lost, the message can be retrieved from the switching center by an action request.

8.3.3 Line Facility Loop Test

The line facilities within the switching center can be tested by connecting the transmit line to the receive line at the switching center data set. This test loop is set up by replacing a pack in the data set with a special loop pack. Test messages can be sent under program control to the transmit channel, looped around to the input or receive channel, and checked by the program. This feature is used to help sectionalize troubles in the link between the control serving test center and the switching center. The test may be requested by an action request from the test center.




Fig. 4—Automatic data channel test facility.

8.3.4 *Automatic Data Channel Test*

The quality of the input signals transmitted from a user's station to the switching center can be determined by the automatic data channel test facility, Fig. 4. Each input line is provided with a bridged connection to a distortion measuring set that tests the quality of input messages. The line to be tested is specified by an action request teletypewriter message. A relay network under program control provides access from any designated line to one of four distortion measuring sets. The test set measures the element transitions within each character and compares them with the theoretical element duration. The highest distortion detected for the duration of the test is indicated to the system. A teletypewriter message reports one of seven ranges of distortion for the line under test. After a valid distortion reading has been obtained, the program releases all connections to the test facility.

An action request teletypewriter message from the control serving test center may request the distortion test on a specific line as part of a fault isolation procedure.

IX. CONCLUSIONS

A great deal of hardware and software has been devoted to implementing the maintenance plan described. In addition to modifying existing No. 1 ESS maintenance programs, approximately 100,000 words of new maintenance programs were written. About 60 percent of the stored program is devoted to maintenance procedures and duplication is used extensively to achieve reliability. A No. 1 ESS ADF office has been in operation since February, 1969.

The performance of the system has been good. As might be expected with any new system, improvements and corrections have been made as weak points in the program and hardware were uncovered. Based on the experience to date and the improvements that have been made, the system is performing as expected and should meet the long term performance and reliability objectives.

X. ACKNOWLEDGMENTS

The maintenance plan is the product of many individuals. The authors wish to acknowledge members of the Maintenance Planning Department headed by Mr. R. W. Downing.

REFERENCES

1. Corbin, J. E., Lehman, H., and Townley, R. C., "No. 1 ESS ADF: Autonomous Data Scanner and Distributor," B.S.T.J., this issue, pp. 2857-2885.
2. Aitcheson, E. J., Ault, C. F., and Spencer, R. G., "No. 1 ESS ADF: Message Store—A Disk Memory System," B.S.T.J., this issue, pp. 2887-2913.
3. Potter, J. L., Strebendt, Mrs. F. B., and Williams, J. R., "No. 1 ESS ADF: Magnetic Tape Subsystem," B.S.T.J., this issue, pp. 2915-2940.
4. Genke, R. M., Harding, P. A., and Staehler, R. E., "No. 1 ESS Call Store," B.S.T.J., 43, No. 5, Part 1 (September, 1964), p. 2147.
5. Downing, R. W., Nowak, J. S., and Tuomenoksa, L. S., "No. 1 ESS Maintenance Plan," B.S.T.J., 43, No. 5, Part 1 (September, 1964), pp. 1988-1993; 2016-2018.
6. Seshu, S., "On an Improved Diagnosis Program," IEEE Transactions on Electronic Computers, EC-14 (February, 1965), pp. 76-79.